



Connect. Accelerate. Outperform.™

Mellanox WinOF VPI Release Notes

Rev 5.10.50000

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT (“PRODUCT(S)”) AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES “AS-IS” WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER'S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies
 350 Oakmead Parkway Suite 100
 Sunnyvale, CA 94085
 U.S.A.
www.mellanox.com
 Tel: (408) 970-3400
 Fax: (408) 970-3403

© Copyright 2015. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, ConnectX®, Connect-IB®, CoolBox®, CORE-Direct®, GPUDirect®, InfiniBridge®, InfiniHost®, InfiniScale®, Kotura®, Kotura logo, Mellanox Connect. Accelerate. Outperform logo, Mellanox Federal Systems®, Mellanox Open Ethernet®, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, Open Ethernet logo, PhyX®, ScalableHPC®, SwitchX®, TestX®, The Generation of Open Ethernet logo, UFM®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

CyPU™, ExtendX™, FabricIT™, FPGADirect™, HPC-X™, Mellanox Care™, Mellanox CloudX™, Mellanox Multi-Host™, Mellanox NEO™, Mellanox Open Ethernet™, Mellanox PeerDirect™, Mellanox Socket Direct™, NVMeDirect™, StPU™, Spectrum™, Switch-IB™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

List of Tables

Table 1:	Beta Features	3
Table 2:	Supported Operating Systems	5
Table 3:	Supported Network Adapter Cards	6
Table 4:	Firmware Versions	6
Table 5:	RoCE Modes Matrix	6
Table 6:	Changes and Fixes in Rev 5.10.50000	7
Table 7:	Generic Issues	11
Table 8:	InfiniBand Issues	14
Table 9:	Ethernet Issues	17
Table 10:	Quality of Service Issues	19
Table 11:	RoCE NVGRE Issues	19
Table 12:	Performance Issues	21
Table 13:	Hyper-V Issues	22
Table 14:	SR-IOV Issues	24
Table 15:	Installation / Upgrade Issues	26
Table 16:	Utilities Issues	30
Table 17:	CIM/WMI Issues	30
Table 18:	ND Issues	31
Table 19:	Change Log History	32
Table 20:	Change Log History	60
Table 21:	API Change Log History	67

Table of Contents

Release Update History	2
Chapter 1 Overview	3
1.1 Main Features in this Release	3
1.2 Beta Features	3
1.3 Unsupported Functionality/Features	4
1.4 WinOF VPI Package Contents	4
1.5 Supported Operating System Versions	5
1.6 Supported Network Adapter Cards	6
1.6.1 Firmware Versions	6
1.7 RoCE Modes Matrix	6
Chapter 2 Changes and Fixes in Rev 5.10.50000 from v4.95.50000	7
Chapter 3 Known Issues	11
3.1 Generic Issues	11
3.2 InfiniBand Issues	14
3.3 Ethernet Issues	17
3.4 Quality of Service Issues	19
3.5 RoCE NVGRE Issues	19
3.6 Performance Issues	21
3.7 Hyper-V Issues	22
3.8 SR-IOV Issues	24
3.9 Installation / Upgrade Issues	26
3.10 Utilities Issues	30
3.11 CIM/WMI Issues	30
3.12 ND Issues	31
Chapter 4 Change Log History	32
Chapter 5 API Change Log History	67

Release Update History

Release	Date	Description
Rev 5.10.50000	October 2015	Updated the following sections for adding support for Windows 10 Client: <ul style="list-style-type: none">• Section 1.5, “Supported Operating System Versions”, on page 5• Section 2, “Changes and Fixes in Rev 5.10.50000 from v4.95.50000”, on page 7
	September 2015	Initial Release

1 Overview

These are the release notes for Mellanox WinOF Rev 5.10.50000 VPI drivers for ConnectX®-3 and ConnectX®-3 Pro adapters.

The driver provides improved performance and additional functionality compared to the Inbox driver provided in Windows Server 2012, Windows Server 2012 R2 and Windows 2008 R2. Therefore, Mellanox strongly recommends updating to its latest driver.

1.1 Main Features in this Release

- **Memory Registration Resources (MTTs) Optimization**
Driver now uses a minimal number of memory registration resources when physical memory contiguity allows it
- **Co-installer**
Allows the installation of counters while installing the driver via INF mechanism
- Added an RDMA test to the package to demonstrate Linux interoperability
- **WPP Tracing**
Mellanox installer now starts an automatic WPP trace session for drivers upon installation.
- **IB Adapter PKey Information in SR-IOV VM**
mlxtool now allows the query of PKeys configured in SRIOV VMs
- **Excessive Pauses Prevention**
Enabled a device watchdog mechanism that prevents the device from sending excessive pauses to the network for any reason
- **MAC Address from Primary Interface**
Windows 2008R2/Windows 7 teaming driver allows selecting the MAC address of the primary
- Added support for driver Teaming in Windows Client 8.1

1.2 Beta Features

Table 1 - Beta Features

Category	Description
WinVerbs	WinVerbs is currently at beta level.
ibdump	ibdump is currently at beta level.
IPoIB	IPv6 support of IPoIB in an SR-IOV guest OS over KVM is at beta level.
	IPoIB teaming support is at beta level and it is supported only on native machine (and not in HyperV or SR-IOV).

1.3 Unsupported Functionality/Features

The following are the unsupported functionalities/features in WinOF Rev 5.10.50000:

- ND over WinVerbs provider
- SRP
- IPv6 over IPoIB Team ports
- VMQ over IPoIB team ports
- Configure IPoIB team through PowerShell
- ConnectX®-2 adapter cards
- IBAL performance tools (ib* ibv*)

1.4 WinOF VPI Package Contents

The Mellanox WinOF Rev 5.10.50000 for Windows package contains the following components:

- Core and ULPs:
 - IB HCA low-level drivers (mlx4)
 - IB Access Layer (IBAL)
 - Ethernet driver (ETH)
 - IP over InfiniBand (IPoIB)
 - NetworkDirect (ND)
- Mellanox Ethernet LBFO driver for Windows Server 2008 R2
- Mellanox IPoIB failover driver
- Utilities:
 - Low level performance tools
- CIM, PowerShell, and WMI support¹

1. Supported in Windows Server 2012 and above, and Windows Client 8.1 and above.

1.5 Supported Operating System Versions

The following describes the supported operating systems and their roles in a virtualization environment.

Table 2 - Supported Operating Systems

Virtualization Mode	Supported Host OS	Supported Guest OS
Native (no-virtualization)	Windows Server 2008 R2 (64 bit only)	N/A
	Windows Server 2012 (64 bit only)	N/A
	Windows Server 2012 R2 (64 bit only)	N/A
	Windows 7 Client (64 bit only)	N/A
	Windows 8.1 Client (64 bit only)	N/A
	Windows 10 Client (64 bit only)	N/A
Hyper-V (non-SR-IOV)	Windows Server 2008 R2 (64 bit only)	<ul style="list-style-type: none"> Windows Server 2008 R2 (64 bit only)
	Windows Server 2012 (64 bit only)	<ul style="list-style-type: none"> Windows Server 2008 R2 (64 bit only) Windows Server 2012 (64 bit only)
	Windows Server 2012 R2 (64 bit only)	<ul style="list-style-type: none"> Windows Server 2008 R2 (64 bit only) Windows Server 2012 (64 bit only) Windows Server 2012 R2 (64 bit only)
SR-IOV Ethernet	Windows Server 2012 R2 (64 bit only)	<ul style="list-style-type: none"> Windows Server 2012 (64 bit only) Windows Server 2012 R2 (64 bit only)
SR-IOV InfiniBand	KVM with MLNX_OFED 2.4 and above	<ul style="list-style-type: none"> Windows Server 2008 R2 (64 bit only) Windows Server 2012 (64 bit only) Windows Server 2012 R2 (64 bit only)

1.6 Supported Network Adapter Cards

Mellanox WinOF Rev 5.10.50000 supports the following Mellanox network adapter cards:

Table 3 - Supported Network Adapter Cards

NICs	Supported Protocol	Supported Link Speed
ConnectX®-3 Pro	InfiniBand (IB)	SDR, DDR, QDR, FDR10, FDR
	Ethernet	10, 40, 50 and 56 Gb/s
ConnectX®-3	InfiniBand (IB)	SDR, DDR, QDR, FDR10, FDR
	Ethernet	10, 40, 50 and 56 Gb/s



The speeds listed in the table above are according to the speeds supported by the device.

1.6.1 Firmware Versions

Mellanox WinOF Rev Rev 5.10.50000 provides the following firmware for Mellanox NICs:

Table 4 - Firmware Versions

NICs	Recommended Firmware Rev.	Additional Firmware Rev. Supported
ConnectX®-3 Pro / ConnectX®-3 Pro EN	Rev 2.35.5100	Rev 2.34.5000
ConnectX®-3 / ConnectX®-3 EN	Rev 2.35.5100	Rev 2.32.5550

1.7 RoCE Modes Matrix

The following is RoCE modes matrix:

Table 5 - RoCE Modes Matrix

Software Stack / Inbox Distribution	RoCE MAC Based (Layer 2) Supported as of Version	RoCE IP Based (Layer 2) Supported as of Version	RoCE v2 (Layer 3) Supported as of Version
Mellanox WinOF	3.2 (Default)	4.80 (Requires additional configuration)	4.70 (Requires additional configuration)
Inbox Windows Server 2012 / Inbox Windows Server 2012 R2	Supported (Default)	Not supported	Not supported

2 Changes and Fixes in Rev 5.10.50000 from v4.95.50000



This package version is Rev 5.10.50000. The package contains the following versions of components:

- Bus, eth, IPoIB and mux drivers version is 5.10.11345.
- The CIM provider version is 5.10.11345.

Table 6 - Changes and Fixes in Rev 5.10.50000

Category	Description	Internal Reference Number
Operating System	Added support for a new Windows Client version - Windows 10 Client	-
General	Fixed an issue where a device state dump used for debug could cause the device to get stuck, requiring driver restart or server reboot to recover	568240
	Fixed an issue where a system with two or more adapters could crash when one of the adapters gets disabled.	532481
	Fixed the case where during very high CPU load on the core that is in charge of transmit notification processing, driver may mistakenly decide the device is hung and initiate a reset to the interface	550016
	Fixed the issue of when a driver fails to start (Mellanox device appears with Yellow bang in the device manager) and user stops the driver, this could cause the system to hang.	492885
	Added third party branding for ATTO. ATTO devices will be shown in the device manager with customized ATTO device names and model numbers.	520073

Table 6 - Changes and Fixes in Rev 5.10.50000

Category	Description	Internal Reference Number
RDMA	Fixed synchronization issue between client and server side in nd_write_lat that could cause the test to hang on start.	559668
	Fixed a problem of when a user is trying to change the RoCE mode using the "Set-MlnxDriverCoreSetting" PowerShell command on a setup with two devices, the mode does not change until the next driver restart.	520406
	Enabled the driver to use a minimal number of memory registration resources when physical memory contiguity allows it.	557943
	Implemented a resource pool to save physically contiguous memory in the driver, which was used by RDMA applications. This way the re-use of this memory will be allowed.	557954
	Improved ND connection establishment time when using firmware v2.34.5000	495620
	Changed NDK and its clients (like SMBDirect) to be disabled by default when PFC is not enabled.	449771
	Added an RDMA test to the package to demonstrate Linux interoperability.	541340
	Enlarged private data limit in ND and NDK.	562879
	Changed default RoCEv2 UDP destination port to match IANA standard.	574918
	Removed all deprecated performance tools.	569889
QoS	Fixed an issue when after OS initiated reset of the interface, default QoS policies were not properly restored.	558513
Virtualization	Fixed the issue of when enabling VMQ after SRIOV has been disabled, VMQ would not work, and the VM would fall back to paravirtualization, impacting performance.	549092
	Improved isolation of SRIOV host from VM driver issues.	549073
	Fixed a race which caused some of the configurations to be badly reinitialized during VM or host transition to VMQ mode. This bug used to cause loss of communication to the affected VM or host.	560789
	Added support in mlxtool to allow the query of PKeys configured in SR-IOV VMs.	565011

Table 6 - Changes and Fixes in Rev 5.10.50000

Category	Description	Internal Reference Number
RoCE	Fixed an issue where adding VLANs would fail because the driver's internal table was not cleaned up correctly upon VLAN removal.	547762
	Fixed a memory leak caused by a race between successful finish of a Management Datagram (MAD) and canceling it.	541447
Debugging	Enabled mlxtool to allow a manual generation of register dumps.	542828
	For performance tuning purposes, debug counters were added to indicate once the driver transmit queue got full.	-
	Added a discard counter for performance analysis.	565011
	Added discard counters per Transport Class. The counters names are in beta and are subject to change.	591908
Ethernet	Fixed the adapter name in command <code>vea_man</code> to not be case sensitive,. For example: <code>vea_man -a ethernet9</code> is now the same as <code>vea_man -a Ethernet9</code> .	538603
	Fixed the issue when after removing the Virtual Ethernet Adapter, some registry keys would remain.	529621
	Enabled a device watchdog mechanism that prevents the device from sending excessive pauses to the network for any reason.	
	Added support for Windows 2008R2/Windows 7 teaming driver to allow selecting the MAC address of the primary interface.	514256
	Added support for driver Teaming in Windows Client 8.1.	507319

Table 6 - Changes and Fixes in Rev 5.10.50000

Category	Description	Internal Reference Number
InfiniBand	Fixed the issue of when Query Path Record (QPR) option is set, a race condition occurs. The race would be between the handling of received packets and the response to the path query to the new destination. This could cause losing packets received from a new source because the path query for it was not yet finished.	536405
	Fixed the Query Path Record list handling to prevent double entering of the same destination. This bug may have caused list corruption which led to unexpected results.	535446
	Fixed the issue of when BSOD may occur when running with two HCAs and using sminfo when no Subnet Manager is available.	492579
	Fixed the issue of when creating virtual IPoIB PKey interfaces with HP cards using part_man.exe utility was not possible.	491585
	Fixed the issue where a new VM creation or its migration in IPoIB could cause the system to crash.	441213
Performance	Fixed an issue where in VMQ mode, not all receive buffers allocated for the VMQ are used, impacting performance.	567513
Installation	Added support for installing counters with co-installer. This allows the installation of counters while installing the driver via the INF mechanism.	549805

3 Known Issues

3.1 Generic Issues

Table 7 - Generic Issues

Reference Number	Issue	Workaround
-	Pinning all the physical memory (used by RDMA operations, such as register memory, pin user memory) on the machine, on Operating Systems prior to Windows Server 2012, may cause the machine to hang.	Avoid pinning the whole machine memory in those Operating Systems.
-	When running applications that use ND or libibmad (such as OpenSM) the system might get to an unstable state when trying to shutdown/restart/hibernate it.	Close all applications that use ND or libibmad before performing shutdown/restart/hibernate.
322721	Activating NC-SI in WinOF v4.90.10541 may cause driver's loading failure when using an older firmware version than 2.30.8000.	Do not enable NC-SI in machines that WinOF v4.90.10541 is installed in.
-	The maximum values returned by the <code>ib_query_ca()</code> function (for example: <code>max_qp</code> , <code>max_mr</code>) are the upper limits of the supported resources by the device. However, it may be impossible to use these maximum values, since the actual number of any resource that can be created may be limited by the machine configuration, the amount of host memory, user permissions, and the amount of resources already in use by other users/processes.	-
-	Running <code>Ntttcp</code> without the <code>"-a X"</code> flag ($X > 1$) in a NIC configured with 10GbE, may cause low bandwidth in TCP single stream.	Run <code>Ntttcp</code> with <code>"-a 8"</code> for best performance
-	Active links disappear after changing the cable connectivity from Ethernet to InfiniBand or vice versa.	Disable and enable the <code>mlx4_bus</code> interface from the Device Manager.

Table 7 - Generic Issues

Reference Number	Issue	Workaround
-	On ConnectX®-2/ConnectX®-3 Ethernet adapter cards, there is a mismatch between the GUID value returned by firmware management tools and that returned by fabric/driver utilities that read the GUID via device firmware (e.g., using <code>ibstat</code>). <code>Mlxburn/flint</code> return <code>0xffff</code> as GUID while the utilities return a value derived from the MAC address. For all driver/firmware/software purposes, the latter value should be used.	Please use the GUID value returned by the fabric/driver utilities (not <code>0xffff</code>).
-	The UI displays the network interface's state incorrectly even after disabling and re-enabling the "disconnected" network interfaces.	To see the correct state and link speed, perform one of the following: <ul style="list-style-type: none"> • Run <code>Get-netadapter Power-shell cmdlet</code> or • Right click on that interface from "Network Connections" and click on status
-	WoL is not supported on Windows Server 2008 R2.	-
-	iSCSI boot over Windows Server 2008 R2 is not supported.	-
410269	Clearing the Mellanox device counters through <code>perfmon</code> does not always work.	Restart the driver
401792	When running in SR-IOV mode and Hyper-V, the same driver version should be used on both the Hypervisor and the virtual machine.	-
-	Working with user space RDMA resources (i.e. ND) in parallel with hardware reset may cause unexpected behavior.	-
491668	In Windows Server 2008 R2 and Windows 7 Client, ConnectX-3 Pro is identified as ConnectX-3. This does not affect anything.	-

Table 7 - Generic Issues

Reference Number	Issue	Workaround
-	On rare occasions, as a result of port configuration change (IB/ETH) the UI may get stuck for up to a few minutes. This effect does not require any user action. The UI returns to its proper functionality after a few minutes.	-
-	The drivers' functionality is limited up to 128 cores.	-
-	According to the hardware architecture, Mellanox NIC devices are presented in the registry as virtual devices and not as physical devices.	Any customer programs or scripts that look into the NIC registry values should not assume it is a physical device.

3.2 InfiniBand Issues

Table 8 - InfiniBand Issues

Reference Number	Issue	Workaround
464449	IPv6 traffic between Hyper-V hosts over IPoIB v-Switch may experience traffic loss.	-
491546	Creating a virtual IPoIB port with non-default PKey isn't supported in the WinOF upgrade. Doing so will indicate that the adapter is unplugged after the upgrade ends.	Remove the virtual port before the upgrade and re-add it after the upgrade.
-	InfiniBand application that was compiled with an SDK version earlier than WinOF v4.90 is not binary compatible.	Recompile InfiniBand application with WinOF v4.90 and above. ND application is backward compatible and older applications over ND do not have to be recompiled.
186806	WinOF 4.40 and above IPoIB driver is not IPv6 compatible with earlier driver versions, including Window 8 Inbox driver. If WinOF 4.50 IPoIB node receives an icmpv6 message from the legacy IPoIB node, the following event will appear in the event log: "IPoIB driver detected a possible incompatibility with Windows 8 inbox IPv6 support due to which there is no IPv6 connectivity".	To enable compatibility mode, add: <ul style="list-style-type: none"> Win8InboxCompatibility-Mode REG_SZ registry key with the value of 1 in the IPoIB interface registry. Note: All IPoIB nodes must use the same mode to enable IPv6 compatibility with earlier driver versions. We recommend upgrading all hosts to the new driver version from http://mellanox.com or use Windows Server 8 compatibility mode
-	Without separate ports for each stream, Win-Sock multiplexes every packet to every subscriber socket and then filters it out.	Use different UDP ports to get higher performance when using multicast packets.
-	A virtual IPoIB interface, created by the <code>part_man</code> utility, reports an Active state when the physical link is in the Initializing state and OpenSM is not running in the subnet	-
-	The "Packets Received Discarded" and "Packets Received Errors" counter may display wrong results.	-

Table 8 - InfiniBand Issues

Reference Number	Issue	Workaround
-	Connection failure on ND tests while machine A have IBAL provider and machine B have MLX4ND provider.	-
-	Hibernate and Sleep are not functional when user-space is using its resources.	-
-	IPoIB does not support: <ul style="list-style-type: none"> • MAC address change • QoS (packet priority) • Connected Mode 	-
-	In an interoperability environment that has both Linux and Windows OSs, the MTU value must be the same, otherwise packets larger than the minimum will not go through. The default MTU for Linux is 2K and for Windows is 4K.	-
-	OpenSM does not run as a service during installation since the SM requires the GUID parameter to decide which port to work on. Setting it on setup causes it to work only on the first port and not the others.	To run OpenSM as a service, assuming the package was installed in the default path, use: <pre>sc create opensm binPath="c:\Program Files\Mellanox\MLNX-_VPI\IB\Tools\opensm.exe"</pre> To start the service run: <pre>sc start opensm</pre>
-	Tools issues: <ul style="list-style-type: none"> • ibportstate does not work on RoCE ports • ibdiagpath may crash on Hyper-V machines 	-
-	If an application which uses InfiniBand runs while the driver is being restarted, a blue-screen or an NMI may occur.	Stop all InfiniBand applications including OpenSM upon driver restart.
330284	If OpenSM is up during driver restart on the same machine, it might stop working, and an error message that reads “Another OpenSM is running” will pop up.	To resume operation, stop OpenSM and restart the driver while OpenSM is down.
-	Sometimes during traffic, the latency of the IPoIB can get higher (this behavior is inconsistent).	Set the following registry in the IPoIB network interface: RecvIntModTime=0 Please note, this action increases the CPU utilization.

Table 8 - InfiniBand Issues

Reference Number	Issue	Workaround
-	No communication between the physical interface and a VM that uses vSwitch created over virtual IPoIB (Pkey), and vice versa.	-
90224	Burning a Firmware image with a "mtu_cap" value other than the default causes the driver load to fail.	Do not change the "mtu_cap" value
383460	Calling <code>ib_query_cq()</code> on a CQ which was created with 1 entry indicates that there are 0 entries in the CQ.	-
439805	IPoIB ports report the configured MTU instead of the effective MTU. These MTUs differ when the corresponding InfiniBand partition is configured with a smaller MTU than the port's configured MTU. In such case, there may be communication failures and/or CPU work overhead on the communicating nodes (even if the peer node is configured correctly).	Make sure the configured adapter MTU (the "Jumbo packet" property in the adapter's advanced settings) is consistent with the corresponding InfiniBand partition MTU. Specifically, the default IPoIB adapter MTU is 4KB and the default OpenSM partition MTU is 2KB. Thus, the default MTU of adapters and OpenSMs is inconsistent. Therefore, the MTU of either all adapters or all OpenSMs (which may be on network switches) on the InfiniBand subnet needs to be changed. Note that the adapter's MTU can be configured through PowerShell.
-	The displayed MAC address in the DHCP server of virtual IPoIB I/F may display a wrong data (FF00.....) although the I/F is still fully functional.	-
489130	DHCP messages that IPoIB send are not fully spec compliant. The spec specifies that the 'chaddr' field must be zeroed, but WinOF IPoIB does not guarantee it.	-

3.3 Ethernet Issues

Table 9 - Ethernet Issues

Reference Number	Issue	Workaround
-	Disabling the “Priority & VLAN tag” in the UI which VLANID is configured, may result in sending packets with the configured VLANID.	Remove the VLANID before disabling the “Priority & VLAN tag”.
206528/ 206945	WakeOnMagicPacket registry key is not added to the registry although WoL is supported by the driver and by the NIC.	-
189704	When the ports of the device are configured as Ethernet only, ibstat/vstat may display wrong information.	-
-	High multicast drop rate on multicast storming.	Use “Multicast traffic” tuning option under the performance tab. For further information, please refer to section “Tunable Performance Parameters” in the User Manual.
-	When there is a stress in TCP connection establishments, some of those connections may fail.	Increase the Ring queue sizes: <ul style="list-style-type: none"> • ReceiveBuffers - controls the receive ring size • TransmitBuffers - controls the transmit ring size
-	The DCB component specifies a default traffic classification that is applied to all egress packets that do not match other classification conditions. In this case, the network adapter assigns the IEEE 802.1p priority level that is associated with the default classification to these egress packets. The default traffic classification has the following attributes: <ul style="list-style-type: none"> • It has a traffic classification condition of type NDIS_QOS_CONDITION_DEFAULT. • It is the first traffic classification defined in the array of NDIS_QOS_CLASSIFICATION_ELEMENT structures. 	-
416248/ 409093	Disabling the ethernet adapter with more than 25 VLANs configured over windows server 2008 R2 \ windows server 7, may result with a non-responding server.	Reduce the number of configured VLANs prior to disabling the ethernet adapter.

Table 9 - Ethernet Issues

Reference Number	Issue	Workaround
410377	Virtual Ethernet Interfaces created by <code>vea_man</code> are not tuned by the automatic performance tuning script.	For optimal performance need to follow the performance tuning guide and apply relevant changes to the VEA interface
-	In Windows Server 2008 R2, devices created by the Teaming driver do not show the correct OEM branding in the Device Manager.	-
326885	Wake on Lan (WoL) cannot be disabled on NICs which supports it.	-

3.4 Quality of Service Issues

Table 10 - Quality of Service Issues

Reference Number	Issue	Workaround
-	Running Quality of Service (QoS) commands without the parameter "-PolicyStore ActiveStore" may cause machines to load without Quality of Service policy.	Store the QoS policy in the ActiveStore

3.5 RoCE NVGRE Issues

Table 11 - RoCE NVGRE Issues

Reference Number	Issue	Workaround
-	RoCE does not support: <ul style="list-style-type: none"> Traffic cannot go through the router. It works in the same subnet only Multicast traffic VLAN Layer 3 feature 	-
327531	In machines with heterogeneous NICs: a NIC which supports RoCE v2, and a NIC which does not support RoCE v2 the following issue might arise: <ul style="list-style-type: none"> ConnectX®-3 Pro is loaded with the transport type RoCE v2 ConnectX®-3 Pro is loaded with the transport type RoCE 	-
-	When NVGRE off-load is enabled, the GRE traffic cannot be accepted as a regular L2 traffic and requires special L2_TUNNELING steering rules. In such case the GRE packets are dropped or directed to promiscuous queue.	-
193619	Using different versions of RoCE in your cluster is not supported.	Use the same RoCE version in all the cluster in the Ethernet ports.
-	RDMA Activity counters do not count during NetworkDirect RoCE traffic.	-
-	GRE traffic steering by inner MAC and by outer MAC simultaneously is currently not supported.	Configure steering or by inner MAC, or by outer MAC.

Table 11 - RoCE NVGRE Issues

Reference Number	Issue	Workaround
-	If VMQ set filter requests are accepted without a GRE flag (i.e. requested steering by outer MAC), the GRE packets do not reach that VMQ.	<p>Set the bus driver registry key- AcceptGREbyOuterMAC_P1/2 per port to accept GRE traffic by outer MAC and to duplicate L2 steering rule to L2_TUNNELING rule for each VMQ set filter request without GRE flag.</p> <p>Note: For regular NVGRE Hyper-V scenarios the value of the registry key below must be set to 0:AcceptGREbyOuterMAC_P1/2</p>

3.6 Performance Issues

Table 12 - Performance Issues

Reference Number	Issue	Workaround
206696	When using WinOF 4.40 or above, low throughput will be seen on 40GbE adapters when QoS is enabled.	Disable QoS when it is not in use: <ul style="list-style-type: none"> Open a PowerShell prompt. Run: <code>Disable-NetAdapterQos -name <Interface Name></code> where <Interface Name> is e.g. "Ethernet 1"
-	<code>perf_tuning</code> is supported only when one of the two NUMA nodes are in use.	-
-	Running performance benchmarks for a short period of time (< 1 sec) may provide bad latency in IPoIB and Ethernet.	Set "Rx Interrupt Moderation Profile" and "Tx Interrupt Moderation Profile", to "Low Latency" to avoid bad latency. Note: This may increase CPU utilization.
-	The driver uses optimal interrupt moderation values for 10 GbE SR-IOV VF scenario. For other scenarios, the optimal values yet to be found.	-
443137, 439897	While running in a Virtual Machine (working with VF) or in a Native Machine, performance counters, if read directly or by using any tool (as Windows Task Manager), may show that no packet has been sent/received. This happens because the driver periodically examines the actual performance counters and caches the results. If the sample rate is too high, the counter values will remain the same.	Reduce the counters sample rate.
587553	Perf Tuning does not allocate RSS cores correctly on all profiles when being used on a platform with more than one Processor Group.	Apply manual tuning via PowerShell.

3.7 Hyper-V Issues

Table 13 - Hyper-V Issues

Reference Number	Issue	Workaround
-	When the vSwitch is detached from the ETH\IPoIB device while the driver is disabled, the device does not reacquire the static IP it had before the attachment of the vSwitch. When the vSwitch is attached to the ETH\IPoIB device while there is no link, it will not receive the device IP when the link is back up.	-
-	After attaching the vSwitch to the ETH\IPoIB device, changing the “Jumbo Packet” registry key on the ETH\IPoIB device does not affect the vSwitch configuration and vice versa. For example, if the user sets the “Jumbo Packet” on the ETH\IPoIB device to X, and the “Jumbo Packet” on the vSwitch to X+Y, X+Y sized packets will be passed from NDIS down to the driver and they will be dropped by it.	Reattach the vSwitch to sync with the value set in the ETH\IPoIB device.
-	Unexpected behavior might occur when running in a virtualized environment and creating two virtual switches bound to each of the ports of a dual port NIC and then using both of them with two vNICs from the same VM.	-
-	In IPoIB when using long Multicast traffic from a Virtual Machine (VM) to an external host there might be up to 0.5% loss in 5% bursts	-
-	Hyper-V is at low bandwidth on LBFO vSwitch, Windows Server 2012	-
-	In Ethernet to achieve better iperf TCP performance between a Linux VM and a Windows VM on different hosts, when using MS MUX over the Ethernet driver, use the non VMQ mode for the VMs.	-
-	After disabling and enabling a port on a guest, a ping to it may be renewed after a minute. The ARP requests sent by Windows are less frequent as the time passes. If the guest port was down for a while, it could take time until Windows decides to send another ARP request to it.	-

Table 13 - Hyper-V Issues

Reference Number	Issue	Workaround
-	<p>When VMQ is enabled after reset, the driver loads all the VMQs that existed before the reset. However, it is not guaranteed that each VMQ will receive the same QP number it had before the reset.</p> <p>This can cause some delay as a result of resetting before connectivity is reestablished. The delay is caused by the time it takes for the ARP table to update after initiating the Gratuitous ARP.</p>	-
-	<p>The IPoIB non-VMQ mode is supported only when the VMQ is enabled according to the registry values.</p>	<p>To use the non-VMQ mode for a VM, change its settings as follow:</p> <ul style="list-style-type: none"> • Press “Settings” on the VM • Go to Network Adapter -> Hardware Acceleration • Un-check the “Enable virtual machine queue”

3.8 SR-IOV Issues

Table 14 - SR-IOV Issues

Reference Number	Issue	Workaround
448196	Working with Windows guest OS over non-Windows SR-IOV hypervisor may result with higher latency compared to Windows Hypervisor.	-
-	When working with OEM adapter cards in SR-IOV mode, the VF is identified as a Mellanox adapter. There is no other effect on the behavior of the VF.	Verify that the FW INI contains the parameter <code>vf_subsystem_id</code> with the OEM encoding in the [HCA] section.
-	An SR-IOV Virtual Machine is loaded in InfiniBand mode when no OpenSM is available in the subnet. The following event appears in the event log: “detected a null port GUID for port <NUMBER>. A Virtual Function device may have a null port GUID if there is no OpenSM instance on its network. Please make sure the network has an active OpenSM and restart the driver.”	Start the OpenSM and restart the driver.
418268	For InfiniBand SR-IOV guest, OpenSM Assigned GUIDs are not supported and may cause unexpected behavior	Work only with Administrator assigned GUIDs.
408734	For InfiniBand SR-IOV guest, ND traffic does not work on an SR-IOV adapter when a Para-Virtualization adapter configured on the same virtual machine and the same subnet as IPoIB.	Disable the Para-Virtualization adapter.
-	In SR-IOV mode over Hyper-V, all ports are in pure Ethernet mode and RDMA is not supported on either port.	-
342421	In SR-IOV mode, enabling SR-IOV in the BIOS may change the interfaces names. If any VSwitch is bounded to an interface whose name was changed, there will not be any way to perform any operation on them.	Unbond all Vswitches from the NIC's interfaces before enabling SR-IOV in the BIOS.
427038	In SR-IOV, 40Gbps interfaces of vSwitch interface on HyperV, and the VMNIC are identified as 10Gbps. Despite the incorrect speed identification, they still achieve 40Gbps performance.	-

Table 14 - SR-IOV Issues

Reference Number	Issue	Workaround
-	Device name in the Device Manager does not show the correct OEM branding for SR-IOV Virtual Function devices without the updated firmware.	-

3.9 Installation / Upgrade Issues

Table 15 - Installation / Upgrade Issues

Reference Number	Issue	Workaround
-	Upgrading the driver while the UI is opened with the “ConnectX NIC device” may cause the installation process to never end.	Close the UI before driver upgrade.
207497	Rebooting the machine while uninstalling WinOF may result in installation failure.	Delete Mellanox components from HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\DIFxApp\Components. The Mellanox components are mlx4eth63, ipoib6x and mlx-4_bus
137859	Canceling the installation process may leave the bus driver in a disable state. The driver appears in a yellow bang containing the following error message: “Windows cannot start this hardware device because its configuration information (in the registry) is incomplete or damaged. (Code 19)”.	Scan for new hardware and reboot the machine.
198537	Downgrade is not supported.	Uninstall the current version and install the older one.
-	Configuration is not restored when replacing a ConnectX®-3 NIC with a ConnectX®-3 Pro NIC located on the same PCI slot	Clean the old network adapter configuration prior to upgrade.
-	Configuration can be restored only in Windows Server 2012 and above	-
-	IPv6 configuration restore is not supported	-

Table 15 - Installation / Upgrade Issues

Reference Number	Issue	Workaround
-	<p>Upon upgrade, the following Registry Key values will be overwritten with the following:</p> <ul style="list-style-type: none"> • *ReceiveBuffers = 512 • *MaxRssProcessors = 8 • *RssBaseProcNumber = 0 • *NumRSSQueues = 8 • *RssMaxProcNumber = 63 • *RssProfile = 1 • DefaultRecvRingProcessor = -1 • TxInterruptProcessor = -1 • TxForwardingProcessor = -1 • RxIntModerationProfile = 1 • TxIntModerationProfile = 1 • RecvCompletionMethod = 1 • SingleStream = 0 • TxRingNum = 8 	-
-	<p>Upon upgrade the following Ethernet Registry Keys will be deleted:</p> <ul style="list-style-type: none"> • SendCompletionMethod • UseRSSForRawIP • UseRSSForUDP 	-
-	<p>Upon upgrade the SendCompletionMethod IPoIB Registry Key value will be modified as follow:</p> <ul style="list-style-type: none"> • SendCompletionMethod = 0 	-
-	<p>Upon upgrade the following IPoIB Registry Keys will be deleted:</p> <ul style="list-style-type: none"> • UseRSSForRawIP • UseRSSForUDP 	-
383584	<p>Uninstalling the driver on Windows Server 2008 R2 with LBFO configuration results in the appearance of a pop-up window requesting to close several running applications.</p>	<p>Choose “Do not close applications”. This action allows the uninstallation of the driver. A Reboot may be required.</p> <p>Rebooting the server before uninstalling the driver when LBFO is configured will eliminate this pop-up completely.</p>
-	<p>Running a downgrade in silent mode is not supported. Upon downgrade the return code will always be 0.</p>	-

Table 15 - Installation / Upgrade Issues

Reference Number	Issue	Workaround
-	Uninstalling the driver after upgrade won't remove the directory %ProgramFiles%\Mellanox	-
-	Uninstalling the driver when multiple of VLANs are configured never ends.	Remove the VLANs before uninstallation.
-	The installation process does not close any applications running in the background, and may cause a BSOD as a result of a stuck cmd.	It is recommended to close all running applications prior to upgrading the driver.
-	Installation/upgrade fails due to PNP failure to copy the driver files to the driver store, and the following text is printed in the event logs: Fault bucket, type 0 Event Name: PnPDriverImportError Response: Not available Attached files: C:\Users\ <user>\appdata\local\temp\dmi151a.tmp.log.xml </user>\appdata\local\temp\dmi151a.tmp.log.xml C:\Program Files\Mellanox\MLNX-_VPI\ETH\mlx4eth63.inf	Reboot the machine and reinstall
403352	Installation/upgrade fails due to failure to stop the WMI service, and the following text is printed in the installation log: "CustomAction StopWMIService returned actual error code 1603"	Kill the WMIPrvSE.exe tasks in the task manager and reinstall.
-	Following the upgrade of Mellanox driver to WinOF-4.60 and above or on servers with no internet access, the first PowerShell command might be stuck for ~2-3 minutes before its completed. According to the following content, this issue is related to .Net framework version or an issue with the internet access: http://www.minasi.com/forum/topic.asp?TOPIC_ID=39253	Run the following script on the server to optimizes loading PowerShell DLLs: \$Env:PATH = [Runtime.InteropServices]::GetRuntimeDirectory() [AppDomain]::CurrentDomain.GetAssemblies() % { \$pt = \$_.Location if (! \$pt) {continue} if (\$cn++) {'' \$na = Split-Path -Leaf \$pt Write-Host -ForegroundColor Yellow "NGENing \$na" ngen install \$pt }

Table 15 - Installation / Upgrade Issues

Reference Number	Issue	Workaround
417380/ 415257	On ConnectX®-3 cards only, when upgrading from Windows Server 2012 R2 Inbox driver to WinOF, the RoCE mode setting in the registry is not properly transferred to the new driver. In case a non-default value was used it will not be configured following the upgrade.	Reconfigure the RoCE Mode setting manually.
427676	Uninstall from the Device Manager is currently unsupported	Driver uninstall can be completed from the Programs and Features window as explained in the Uninstalling Mellanox WinOF Driver section in the User Manual.
-	WinOF Inbox driver does not support upgrade. When installing WinOF v4.40 and above on a Windows Server 2012 and above machine, the Inbox driver is uninstalled prior to starting the new installation and any previous configurations is lost. The Inbox driver will be reinstalled automatically when the new driver is uninstalled.	-
-	Firmware upgrade may fail during installation if there was a prior firmware upgrade on the machine without a reboot after it. A firmware upgrade failure does not fail the whole installation.	Upgrade the firmware manually.
-	Driver installation requires deletion of the <code>mlx-4_bus.sys</code> file in Windows Server 2008 R2 when using the PXE package.	Delete the <code>mlx4_bus.sys</code> file and reboot the machine to install the driver.
492398	If there are disabled network interfaces or a disabled Mellanox bus driver, they will be enabled after the WinOF upgrade.	-
-	The upgrade from WinOF v4.58 to WinOF v4.90 saves only the configuration of the Ethernet interface, and does not save the configuration of the IPoIB interface.	-
549805	Uninstalling mlx4 bus device from the Device Manager results in the removal of performance counters for all WinOF devices.	-

3.10 Utilities Issues

Table 16 - Utilities Issues

	Issue	Workaround
-	<code>ibdump</code> may encounter packet drops upon a burst of more than 4096 (or $2^{\text{max-burst}}$) packets.	-
-	Packets loss is not reported by <code>ibdump</code> .	-
-	Running <code>ibdump</code> on a RoCE Ethernet port may decrease the functional bandwidth due to the overhead of creating extra copy for each packet. This may lead to packet drops on the link.	Verify Ethernet flow control is enable to ensure a lossless link
-	Pcap file issues for RoCE IP Based: <ul style="list-style-type: none"> The packets 'capture-time' field is wrong (may be 0 or negative numbers). For every captured packet, an additional 0 size frame is added. This appears in Wireshark as a 'malformed Ethernet packet'. 	-
-	Sniffing over IB ports is currently not supported	-
199079	When the tracer tool traces many events, it may consume a large amount of memory (up to several GB RAM).	Use the button to set maximum number of the displayed lines at a maximum buffer size.

3.11 CIM/WMI Issues

Table 17 - CIM/WMI Issues

Reference Number	Issue	Workaround
-	Running Microsoft CIM cmdlets operations and their derived classes on classes <code>MLNX_NetAdapterSettingData</code> and <code>MLNX_NetAdapterRoceSettingData</code> is not supported. Calling those commands may cause the debugger, if connected to the machine, to assert.	Use <code>DriverCoreSettings</code> instead.
408230	For PCI Gen3, <code>PcieLinkSpeed</code> is reported as "Unknown" when running <code>Get-NetAdapterHardwareInfo</code> Powershell cmdlet	-
-	WMI does not work due to lack of permissions.	Change the execution policy. Run: <code>Set-ExecutionPolicy AllSigned</code>

Table 17 - CIM/WMI Issues

Reference Number	Issue	Workaround
433986	The information that is printed in the cmdlets get-net-adaptersriov and Get-MlnxPCIDeviceSetting is inconsistent.	Use only the cmdlet Get-MlnxPCIDeviceSetting.

3.12 ND Issues

Table 18 - ND Issues

Reference Number	Issue	Workaround
432674	There is an interoperability problem between NDK and ND in RDMA operations during the write and read operations. However, the send operation resumes working. This happens since ND uses remote token in a network order (big endian) while NDK uses remote token in CPU order (little endian). Therefore, an inconsistency is caused between ND and NDK in RDMA operations.	An ND application that works with NDK using RDMA operations must handle this issue by changing the remote token to the appropriate byte order before sending it to NDK.
-	When working with the default NDv1 and NDv2 providers, the following error message might be displayed: 0xC0000120 NT_STATUS_CANCELLED This error does not affect any functionality and can be safely ignored.	-
-	Changing the default ND providers may cause random errors, such as: 0xC0000238 (NT_STATUS_ADDRESS_ALREADY_ASSOCIATED) on Connect() or with 0xC0000236 (NT_STATUS_CONNECTION_REFUSED) on Accept(). These errors can be safely ignored.	-

4 Change Log History

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
4.95.50000	InfiniBand	Fixed BSOD on next driver restart when running the sminfo tool and SM is not running	492579	This package version is Rev 4.90.51000. The package contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.95.10777 • The CIM provider version is 4.95.10777
		Fixed instability in IPoIB driver when polling mode is enabled	521205	
		Fixed the issue when live-migrated VM with virtual NIC over IPoIB physical interface loses its IP address and comes up with APIPA address (169.254.x.x)	439359	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
	Ethernet	Fixed BSoD after the receive buffer's size changes in VMQ mode	500228/522073	
		ECN is now configurable via PowerShell	495158	
		Fixed Powershell setting of RoCE mode when the machine has more than one Mellanox device	520406	
		Fixed reporting of NVGRE capabilities to the OS	535203	
		Added a new mode that ignores FCS warnings and enables the Ethernet packets to be received by the NIC	-	
		<ul style="list-style-type: none"> • Added the option of setting the MAC address of 2008R2 teaming driver to be taken from the primary interface • Added the option of taking the MAC Address used for teaming from one of the interfaces without manipulation 	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
	Changes in UI	Added the option of configuring the team MAC address to be the same as the primary adapter MAC address	514989	
	Troubleshooting	Added an autologger session that dumps WPP traces to file to enable easier analysis of issues	-	
		Added logging of performance counters and CPU power saving state to system snapshot tool	-	
	RDMA	Fixed handling of send request with inline data larger than supported	443355	
		Added the option to allow RDMA programmers to create pre-allocated pools of ND resources to reduce resource creation time	-	
	Low Level Driver	Fixed the case in which the system rarely gets unstable after driver restart	492885	
	Infrastructure	Updated customization for OEM cards	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
4.90.50000	Generic	Fixed driver instability when handling many RDMA connection requests in parallel	461854	This package version is Rev 4.90.50000: The package contains the following versions of components. <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.90.10714 • The CIM provider version is 4.90.10714
		Added to MLNX_System_Snapshot Mellanox specific counters and data from Get-Mlnx* Cmdlets	467529	
	Resiliency	Reset Flow improvements: <ul style="list-style-type: none"> • Resolved race condition when reset is initiated by more than one source • Reset initiated on one port does not cause reset of the other port 	400887	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
	Ethernet	ETS is now configurable through DSCP values. For further details, please refer to WinOF User Manual, "Differentiated Services Code Point (DSCP)" section	434105	
		Fixed the issue of when creating a Virtual Ethernet Adapter interface and removing it immediately a Blue Screen may appear	456279	
		Fixed duplicated values of Receive Completion Method in Advanced Properties driver dialog on Windows Server 2012 R2	443273	
		Performance Improvement: Reduced memory access time for Receive descriptors	-	
		VM Scalability: More efficient handling of VMQ control path in HyperV	-	
		Reduced the amount of kernel memory used for each Ethernet interface by the driver	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Virtual Machine traffic on the default queue now uses a single CPU core as required by Microsoft. This applies both to SRIOV and VMQ	441581	
	InfiniBand	Updated IBAL interface version. In order for the applications that use the IBAL interface to work with WinOF Rev 4.90.50000, they must be recompiled with the new SDK	-	
		Added support for SM change event	435564	
		Fixed propagation of error code when <code>ib_join_mcast()</code> fails	448028	
		Fixed connectivity problems when using PKeys from the same partition with different membership types	417753	
		Fixed VM reset after printing the message "mlx-4_core 0000:05:00.0: unparavirt command: OTHER (0x3a) accepted from slave:3" in SR-IOV InfiniBand VM over non-windows hypervisor	422598	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
	RoCE	In RoCE v2, added the option of determining the source port field of the UDP header by the application	-	
	NDK	Improved CPU utilization by changing ndkgetremotetokenfrommr() to return value in network byte order	-	
	Performance	Fixed the UI crash when working with a single port	427484	
		Increased the accuracy of the run time duration parameter of ND Performance tests even when sending large message	-	
		Fixed Maximum value for Thread-Poll parameter to be 200,000, instead of the previous state when it could not be set above 20,000 due to a bug	481291	
	IPoIB	Fixed stability issues	-	
		Fixed displaying of IPoIB default turning option	428601	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Fixed temporary network connectivity issues while migrating VMs or modifying VMQ configuration for VMs that uses IPoIB with VMQ	417687	
		Fixed the part_man tool to use the actual default p_key instead of 0xffff	417858	
		Fixed NIC reset when attaching to a multicast group fails	423435	
		Fixed duplicated values of Receive Completion Method in UI on Windows Server 2012 R2	-	
		Added support for multiple PKey interfaces in IPoIB	-	
		Added support for teaming of IPoIB interfaces to allow failover	443273	
		Added IPoIB adapters teaming support (beta level)	-	
		Added sending of gratuitous ARP in IPoIB interface when the MAC address is changed	408388	
		Reduced memory footprints of IPoIB interfaces	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Reduced the multiple number of path record queries to one when old query information exists	466336	
		Improved completion memory access speed	440018	
		Changed default VMQ/VPport affinity to use first RSS CPU	442549	
		Multiple PKey support is now at GA level. The <code>part_man</code> tool allows the creation of up to 64 vIPoIB interfaces (32 per port)	-	
		Added a warning to the event log if the port MTU is higher than the reported MTU by the SM.	-	
	Installation	Fixed CIM failure after installation in maintenance mode	423206	
		Fixed loading of old driver after driver upgrade that requires system reboot to complete the process	422812	
		Fixed RoCE disable by default after installation of WinOF in Windows 8.1 Client	454020	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
	ND	Fixed seg fault when executing ND application with no device installed or when a wrong device identifier is used	431113	
		Fixed wrong reported value of supported number of SGE in 32 bit DLLs	425841	
		Increased the number of supported SGEs in 32 bit DLLs to 2	425841	
	NVGRE	Fixed restoration of NVGRE configuration after NIC reset	442478	
	Changes in UI	Replaced the terms “LBFO” and “Bundle” with “Teaming” and “team” respectively.	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
	CIM/WMI	Added support to query/set/enable/disable ECN		
		Added support to query Dropless-Mode state		
		Fixed the issue when using the PowerShell command Get-Mlnx-FirmwareIdentity on a system with multiple NICs/HCAs while one of the devices is disabled and the command fails		

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
4.80.50000	Installation/ Upgrade	Added check for administrator privileges during installation	391704	<p>This package version is 4.80.50000. The package contains the following versions of components.</p> <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.80.10388 • The CIM provider version is 4.80.10388
		Added support for installation in silent mode without execution of perf_tune	397946	
		Fixed installation stuck when Remote Desktop Session Host Windows Installer RDS compatibility is enabled	371541	
	Generic	Changed Reset Flow (+SR-IOV)-enabled only if no user space application is running and depends on the registry key: AllowResetOnError setting)	370536	
		Changed the number of supported QPs in a multicast group from hard coded value to firmware capabilities dependent	401850	
		Fixed driver load failure in machines with 1 TB memory and above	407556	
		Fixed memory leak on the Virtual Machine in SR-IOV when resetting the Virtual Machine of associated VFs	373144	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
	IPoIB	Added multiple P_Key support (beta level)	391240	
		Added IPoIB SR-IOV over KVM and ESX Hypervisors (for both full and partialmembership)	-	
		Added support for LID change event	-	
		Added enhancements in <code>part_man</code> for the multiple Pkey support	-	
		Changed IPv6 “all dhcp servers” mcast to be persistent	-	
		Fixed rare cases of driver hang following a Subnet Manager failover event	-	
		Fixed stability issues	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
	Ethernet	Added RSS in UDP (enabled by default)	-	
		Added 56 GbE (Please refer to the Infiniband Switch User Guide for further details)	-	
		Changed DSCP configuration to be per port instead of global	394703	
		Network Direct: Fixed race in NDK between handling of incoming connection and destruction of a listener	-	
		Network Direct: Fixed race between NDK object creation and usage	-	
		Improved TCB (Transmission Control Block) management on send	389974	
		Improved transmit and receive in multi stream scenarios	-	
		Enabled hardware checksum offload for non TCP/UDP traffic with ConnectX®-3 Pro	394977	
		Improved stability when handling OIDs during driver reset	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Fixed performance tuning for 1GbE link	-	
		Fixed possible reset of driver during migration of large number of VMs at the same time	401655	
		Fixed stability issues	-	
	RoCE	Added RoCE IP based	391238	
	ND	Fixed wrong return value in IND2-Adapter::QueryAddressList	-	
	InfiniBand	Added non-default PKey in VM	-	
	Performance	Optimized interrupt moderation values in SR-IOV VF mode for IPoIB	-	
		Improved perf_tuning detection for the first port	-	
		Improved performance in packet forwarding scenarios	-	
		Decreased dropped packets rate for Ethernet significantly	414872	
		Changed default perf_tuning scenario to be "Balanced configuration"	396981	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Various performance improvements	-	
	WMI/CIM	Added ability to read active RoCE configuration from hardware	400598	
		Added support for RoCE IP Based	390573	
4.70.50050	IPoIB	Fixed SM fail-over causing the driver to hang	-	This package version is 4.70.50050. The package contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.70.10143 • The CIM provider version is 4.70.10143
4.70.50040	Generic	Optimized handling of “affinity change” on <code>OID_RECEIVE_FILTER_QUEUE_PARAMETERS</code>	-	This package version is 4.70.50040. The package contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.70.10141 • The CIM provider version is 4.70.10141
		Added the ability to control the number of retries and timeout to check the device health before performing reset	-	
	Ethernet	Fixed missing pause response by sender when using DSCP/untag priority tag mode with ETS enabled	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
4.70.50000	Installation/Upgrade	Fixed removal of virtual IPoIB ports in uninstallation	-	This package version is 4.70.50000. The package contains mixed versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.70.10126. • The CIM provider version is 4.70.10130.
		All user-space binaries are now signed	-	
		Fixed restoration process of DNS servers during upgrade	-	
		Fixed popping windows during installation/upgrade	-	
		Fixed missing 32 bit files in the catalog files	-	
	Generic	Changed Ethernet and IPoIB event log messages to be more clear	-	
		Ported SDK project to Visual Studio 2013.	-	
		Fixed an issue which caused Mellanox miniport devices to be listed in “Devices and Printers”	-	
		Fixed Ethernet and IPoIB deadlock in power state change during shutdown/reboot	-	
		Fixed stability issues	-	
	IPoIB	Added support for IPoIB SR-IOV Virtual Function (VF) over KVM Hypervisor (Beta level)	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Added support for non-default pkey, as queried from OpenSM, on IPoIB SR-IOV VF over KVM.	-	
		Added IPoIB QoS proprietary counters, diagnostics and traffic for monitoring, using Windows' perfmon utility	-	
		Fixed <code>part_man</code> exit with return value 0 in case of error	-	
	Ethernet	Added support for Ethernet SR-IOV over Windows Hyper-V Hypervisor (over Windows Server 2012 R2)*	-	* Requires firmware v2.30.8000 and above
		Added Virtual Ethernet Adapter support which enables using SMB Direct and HyperV (VMQ and NVGRE (over ConnectX®-3 Pro)) on the same port**	-	** Requires firmware v2.31.5050 and above.
		Added lossless TCP buffer management when no receive WQE are available	-	
	RoCE	Added ConnectX®-3 Pro support for RoCEv2	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Changed the transport name in vstat and ibstat to be RoCE v2.0	-	
		Fixed ibstat behavior on devices with RoCE enabled	-	
		Fixed releasing of RDMA resources and reacquire them on power down and up.	-	
		Fixed RDMA Activity counters which didn't increase for ND traffic	-	
	ND	Fixed hard-coded limitation of 4 SGEs	-	
	InfiniBand	Fixed vstat printing of counters for Ethernet ports	-	
		Fixed crash when calling <code>ib_join_mcast()</code> with <code>timeout_ms = 0</code>	330740	
Performance	Improved <code>perf_tuning</code> setting in single CPU machines to avoid TX collision	-		
4.61	Installation/Upgrade	Fixed an issue preventing JumboPackets registry key to be restored correctly	-	WinOF VPI version 4.61 was released as an intermediate release.

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Ensured that uninstillation of Mellanox package in Virtual Machine leaves the system clean	-	
	Generic	Improved information in event log when a bad cable is detected	-	
		Improved resiliency on error flow in Ethernet, IPoIB and bus drivers	-	
		Fixed an issue which caused Mellanox devices to be listed in “Devices and Printers” and had “Safe Removal” UI	-	
	Performance	Added support OF IPv6 to all nd_*_* tests	-	
		Enabled optimal interrupt moderation values in SR-IOV VF mode	-	
		Stopped using NdisQueryNet-BufferPhysical-Count to improve CPU utilization	-	
	IPoIB	Enabled searching for IBAT routes based on dest only instead of src,dest and added a mechanism preventing memory growth in IBAT entries	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Allowed any number of RSS processors, not only a power of 2	-	
		Ensured SR-IOV mode is not enabled for IPoIB ports, which resulted in confusing message in event log	-	
		Fixed error statistics collection which could cause false error report	-	
		Fixed a connectivity problem between Hyper-V VMs on the same host	-	
		Fixed loopback issues in the virtualization environment	-	
		Fixed stability issues	-	
	Ethernet	Added support for “unknown” link state indication	-	
		Added support for DMA checks by driver verifier on SR-IOV Virtual Function	-	
		Added support for NVGRE over LBFO Team	-	
		Improved performance of handling change receive ring affinity request	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		In SR-IOV mode, improved resiliency to driver failures in the Virtual Machine which could result in driver load failure in VM	-	
		In SR-IOV mode, improved resilience in VF to PF communication	-	
		Improved structure of INF file for SR-IOV Physical and Virtual Functions	-	
		Fixed an issue that prevented receiving ARP traffic in NVGRE mode	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
4.60.17718	Hyper-V	Fixed NIC reset when moving IPoIB interface in a VM from non-VMQ to VMQ or from VMQ to non-VMQ	325607	<p>This package version is 4.60.17718. The package contains the following versions of components:</p> <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.60.17718. • The CIM provider version is 4.60.17718. <p>*** Requires firmware v2.30.8000 and above.</p>
	Installation/Upgrade	Enabled configuration changes saving upon Inbox and previous releases upgrade	-	
		Enabled CIM installation as a standalone package	-	
		Fixed an issue occurred when uninstalling and reinstalling the driver. The ConnectX-3 Pro Ethernet device was displayed in the Device Manager with a yellow bang (!).	-	
		Fixed an issues enabling the package's execution in modify mode resulting in driver being disabled	-	
	Generic	Added support for a new report for bad cables ***	-	
		Fixed random parsing failures of string registry entries	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Fixed compilation failure of “Hello_world” in the SDK	-	
		Fixed the return value of <code>ib_query_ca()</code> if failed to allocate resources for operation	-	
	Performance	Added support to IPv6-to-all <code>nd_*_*</code> tests	-	
		Fixed CPU utilization report in <code>nd_*_*</code> tests	-	
		Fixed correct bandwidth peak results in <code>ibv_send_bw</code> with UD QP	-	
		Fixed sync problems of bidirectional mode in <code>ibv_read_bw/ibv_write_bw</code>	-	
		Fixed an issue reporting incorrect adapter type in performance tuning log file	-	
	RoCE	Fixed RoCE mode parsing	-	
	ND	Added the ability to rearm a CQ in the kernel	-	
		Added the ability to handle LID changes	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Changed connection timeout behavior. Added the STATUS_CONNECTION_REFUSED return value upon connection timeout.	-	
		Fixed missing completions when working with Completion Queue with single entry	-	
	IPoIB	Added the ability to handle LID changes	-	
		Added support for iSCSI boot over IPoIB	-	
		Fixed unexpected behavior upon QP asynchronous event	-	
		Fixed bad completions of VMQ and NonVMQ modes in IPoIB	-	
		Fixed a failure occurred when setting the IPoIB adapter value to "SA Query Timeout"	-	
		Fixed propagation of the physical link disconnection to virtual (part_man) interface	-	
		Fixed BSOD caused by calling <code>ib_join_mcast()</code> with <code>timeout_ms = 0</code>	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Performance improvements in latency	-	
	Ethernet	Added DSCP support over IPv4 ^a	-	
		Added traffic profile	-	
		Added IRQ dynamic moderation	-	
		Modified the CQ size to prevent CQ overrun	-	
		Changed the report link speed zero in case of disconnected network adapter	-	
		LBFO: Fixed port channel teaming with CISCO switch and Fabric Extenders traffic loose in Windows Server 2008 R2	-	
		Fixed an issue related to packets sent with corrupted VLAN header when they were meant to be untagged	-	
		Fixed unexpected behavior upon QP asynchronous event	-	
		Fixed the ability to disable Wake on Lan (WoL) on NICs which supports it.	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Stability fixes	-	
		Performance improvements	-	
	WMI/CIM	Added ControlledBy association to IBPort	-	
		Fixed ConformsToProfiles association for SoftwareIdentity and DriverIdentity	-	
		Fixed execution of all tests which were running when executing Diagnostic tests on one instance	-	
		Fixed a failure occurred when running MLNX_Card	-	
		Fixed the printing of diagnostics log	-	
		Fixed an issue preventing from get-event to show information after disabling the PCI device	-	

Table 19 - Change Log History

Release	Category	Description	Reference Number	Notes
		Removed support for the following configuration: <ul style="list-style-type: none"> • ModeFlags • SingleMsixNum • MultiMsixNum • SingleEqNum • MultiEqNum • MaxContQuant • SlaveNum • DebugLevel • DebugFlags • UsePrio • NumFcExch • EnableQoS • BlockMcastLoop-Back • InterruptFrom-FirstPacket • ProbeVf 	-	
4.60.17738	IPoIB	Fixed using CQ after VMQ is closed	-	This package version is 4.60.17738. The package contains the following versions of components: <ul style="list-style-type: none"> • Bus and eth driver version 4.60.17718. • The CIM provider version is 4.60.17718. • The mux driver version is 4.60.17729. • The IPoIB driver version is 4.60.17736.
		Fixed bad completion of VMQ QP that was caused by malformed WR	-	
	Ethernet	LBFO: Fixed the team's MAC address uniqueness in the subnet of the team in Windows Server 2008 R2	-	

Table 20 - Change Log History

Release	Category	Description
4.55	Generic	<ul style="list-style-type: none"> Added support for Windows Server 2012 R2 Operating System Added the <code>ParentBusPath</code> option to each port registry key Added a new hardware ID for ConnectX®-3 Pro NICs The QP numbers allocation is now round-robin manner <code>RecvCompletionMethod</code> as <code>Interrupt</code> is no longer supported Removed the <code>LsoV1IPv4</code> from the registry/UI Removed from the bus driver configuration the '<code>Non-DMA</code>' option Removed the <code>TXRingNum</code> option from the UI
	NVGRE	<ul style="list-style-type: none"> Added NVGRE hardware off-load support (for ConnectX®-3 Pro cards only) Added to the UI the <code>*EncapsulatedPacketTaskOffload</code> option when using ConnectX®-3 Pro NICs
	Performance	<ul style="list-style-type: none"> Added the <code>nd_send_bw</code> and <code>nd_send_lat</code> ND benchmarking tools Fixed <code>nd_*_bw</code> to achieve better performance (memory buffer alignment) and consistent results
	Ethernet	<ul style="list-style-type: none"> Fixed the issue preventing messages to be sent in VLAN 0 when using many VMQ rings Added IP-IP checksum off-load support Added Ports TX arbitration/Bandwidth allocation per port The following ND providers, <code>MLX4ND</code> and <code>MLX4ND2</code> are installed by default Fixed setting the correct SL in UD traffic over RoCE
	InfiniBand	<ul style="list-style-type: none"> IPoIB performance improvements Fixed a <code>part_man</code> issue related to wrong statistics over virtual <code>partman</code> interfaces
	RoCE	<ul style="list-style-type: none"> Enabled <code>roce_mode</code> value overwrite in case it exists during installation Fixed in <code>ibv_devinfo</code> the display of correct transport RoCE mode Added Sniffer for RoCE packets The used RoCE mode set upon driver load is printed into event log message

Table 20 - Change Log History

Release	Category	Description
4.40	Generic	<ul style="list-style-type: none"> • Added a notification in the event log in case SMB is not supported in ConnectX®-2 firmware • Added the trace tool for WPP tracing • Added copyright to the SDK files • Added WMI/Powershell support • Fixed an issue causing the setup to fail upon <code>perf_tuning</code> failure during the installation. An error message will be printed in the installation log upon <code>perf_tuning</code> failure. • Removed port setting registry key during uninstall • Fixed an issue with the Mellanox adapter being shown on the USB removal menu, which caused the removal of the Mellanox adapter once removing the USB.
	Performance	<ul style="list-style-type: none"> • Set 512 RX buffers by default • Removed <code>TXRingNum</code> • Changed the <code>perf_tuning</code> setting to achieve a better performance tuning • Added the <code>nd_write_bw/nd_write_lat</code> and <code>nd_read_bw/nd_read_lat</code> tools • Fixed the <code>perf_tuning</code> indication of the last chosen tuning scenarios • Fixed a crash in the <code>ib_send_lat/bw</code> utilities caused when the port link was down • Fixed the “Restore to defaults” option in the <code>perf_tuning</code> tool. Now the default values are being restored
	Ethernet	<ul style="list-style-type: none"> • Added Transmit Side Scaling (TSS) • Added Ethernet QoS proprietary counters, diagnostics and traffic for monitoring, using Windows’ <code>perfmon</code> utility • Added to the MTU size the IP header size (1500 ->1514, 9600->9614). Thus the minimum Jumbo frame size is 614. • Interrupt moderation supports the following profiles: <ul style="list-style-type: none"> • Low Latency • Moderate • Aggressive <p>In addition to old values that are not supported anymore.</p> • Made <code>mlx4_bus</code> and Ethernet devices removable • Network Direct: Added support for NDv2 • Network Direct: Set the default ND provide value to <code>mlx4nd2</code> • Fixed WoL support on NIC with a single port • Fixed the default RoCE configuration on NICs with a single ports • Fixed the values for the MTU and rate of the CM-REQ • Fixed miniport reset on sending scenarios • Removed the QoS attributes when disabling QoS

Table 20 - Change Log History

Release	Category	Description
		<ul style="list-style-type: none"> • Enabled MaxRssProcessors support of the following values: 1, 2, 4, 8, 16, 32, 64 • Network Direct: Fixed a crash occurred when more than 4 SGEs elements were used in an ND write operation • Network Direct: Fixed the swap of InboundReadLimit and OutboundReadLimit when creating an EndPoint and in Connector::GetConnectionData • Network Direct: Fixed disallowing creation of EndPoint with zero attributes in the Receive Queue • Network Direct: Removed the option of NDK registration failure requiring a reboot of the machine to register it again • Network Direct: Fixed a failure when creating an EndPoint with zero attributes in the Receive Queue • Network Direct: Added the option of sensing the incoming Read messages according to the device capabilities when creating an EndPoint limit • Network Direct: Fixed a failure of ND connectivity between VMs on the same host • Added Transmit Side Scaling (TSS) • Added Ethernet QoS proprietary counters, diagnostics and traffic for monitoring, using Windows' perfmon utility • Added to the MTU size the IP header size (1500 ->1514, 9600->9614). Thus the minimum Jumbo frame size is 614. • Interrupt moderation supports the following profiles: <ul style="list-style-type: none"> • Low Latency • Moderate • Aggressive <p style="margin-left: 20px;">In addition to old values that are not supported anymore.</p> • Made mlx4_bus and Ethernet devices removable

Table 20 - Change Log History

Release	Category	Description
	Ethernet	<ul style="list-style-type: none"> • Network Direct: Added support for NDv2 • Network Direct: Set the default ND provide value to mlx4nd2 • Fixed WoL support on NIC with a single port • Fixed the default RoCE configuration on NICs with a single ports • Fixed the values for the MTU and rate of the CM-REQ • Fixed miniport reset on sending scenarios • Removed the QoS attributes when disabling QoS • Enabled MaxRssProcessirs support of the following values: 1, 2, 4, 8, 16, 32, 64 • Network Direct: Fixed a crash occurred when more than 4 SGEs elements were used in an ND write operation • Network Direct: Fixed the swap of InboundReadLimit and OutboundReadLimit when creating an EndPoint and in Connector::GetConnectionData • Network Direct: Fixed disallowing creation of EndPoint with zero attributes in the Receive Queue • Network Direct: Removed the option of NDK registration failure requiring a reboot of the machine to register it again • Network Direct: Added the option of sensing the incoming Read messages according to the device capabilities when creating an EndPoint limit • Network Direct: Fixed a failure of ND connectivity between VMs on the same host
	InfiniBand	<ul style="list-style-type: none"> • On rare occasions, depends on the GUID assignment, the IPoIB MAC address can be assigned with a multicast MAC (the least significant bit of the most significant address octet is set to 1). In that case, all of the traffic over the IPoIB I/F is dropped. If you experience this issue, please contact Mellanox support. • Added <code>active_mtu</code> field to <code>struct ib_port_attr_t</code> • Added the option of <code>vstat</code> displaying the <code>active_mtu</code> of the ports • Allowed registration of a large Memory Region which is split to many segments • Fixed a bluescreen issue that occurred when disabling the interface after a TX stress over the VMQ • Fixed a failure of MPI/ND over InfiniBand • Added the option of <code>ibv_devinfo</code> displaying the correct MTU value after it was changed • Added the option of <code>part_man</code> printing the adapter name when the Port GUID is set to zero. • Added the option of <code>part_man</code> printing the leading zeroes of port GUID

Table 20 - Change Log History

Release	Category	Description
	Installation/ Upgrade	<ul style="list-style-type: none"> Prevented displaying a message to upgrade the firmware for OEM NICs if it has the latest firmware version Removed <code>portsetting</code> registry key during uninstallation
4.3- This version was released as an intermediate release.	Generic	<ul style="list-style-type: none"> Added support for a new provider called MLX4ND, which supports both NDv1 and NDv2 interfaces
	Performance	<ul style="list-style-type: none"> Enabled performance tuning running according to the operating systems that are running over it. <p>The keywords added to the registry in NDIS support Windows Server 2012 are:</p> <ul style="list-style-type: none"> RssMaxProcNumber NumRSSQueues RSSProfile <p>The rest of the keywords are added in all versions of NDIS. This change is based on: http://msdn.microsoft.com/en-us/library/windows/hardware/ff570864(v=vs.85).aspx</p>
	Ethernet	<ul style="list-style-type: none"> RoCE MTU value is no longer set to 1024 by default. All options stay as they are and can only be chosen if they were selected explicitly in the UI/registry. <p>The current default state is as follows: The value is now derived from the MTU (or MaxFramSize, or Jumbo Packets value) and they are all aliases for the same value). The value is aligned to 256,512,1024,2048 in a way that it will be rounded down to the nearest power of two of the ETH MTU.</p>
	InfiniBand	<ul style="list-style-type: none"> Added <code>ibdiagnet</code> utility support

Table 20 - Change Log History

Release	Category	Description
4.2	Generic	<ul style="list-style-type: none"> • Modified RSS cores and changed VMQ affinity on the fly • Fixed restart issue when there are not enough MSI-X vectors for each machine core • Added support for K-GROUPS processors (more than 64 processors support) to allow assignment of MSI-X affinity for multiple processor groups. • Set an adequate number of MTTs to map all physical memory • Allocated firmware and ICM memory in chunks of non-paged memory instead of using contiguous physical memory. • Fixed RSS indirection table mapping building when there are less RX rings than RSS cores. • Fixed a bug, preventing standard work with BAR value more than 4GB. • Fixed memory leaks • Fixed error flows causing a Bluescreen in driver startup/unload • Fixed a Bluescreen occurrence upon shutdown due to leak in active resources • Changed device names in device manager and their hardware IDs. The changes were made to distinguish between ConnectX®-2 and ConnectX®-3: <ul style="list-style-type: none"> • for ConnectX-2: MLX4\ConnectX-2_Eth and IBA\ConnectX-2_IPoIB • for ConnectX-3: MLX4\ConnectX-3_Eth and IBA\ConnectX-3_IPoIB • Set QoS settings only for ConnectX-3. Changing the hardware ID, forces the OS to install new device and re-build the registry keys. • Added an event log to indicate driver failure upon start if there are two HCA burned with the same GUID.
		<ul style="list-style-type: none"> • Added firmware upgrade support as part of the setup process. The setup burns the new firmware only on Mellanox cards. Firmware burning failure does not prevent the driver's installation, therefore, it will show a warning. In this case, it is recommended to update the firmware manually. • Enabled configuration of TxRingNum registry key from the UI
	Generic	<ul style="list-style-type: none"> • Improved the "Port Protocol" dialog • Added Registry key documentation to the setup package

Table 20 - Change Log History

Release	Category	Description
	Performance	<ul style="list-style-type: none"> • Optimized code performance • Increased send parallelism • Memory used in receive flow is now allocated with the same affinity of the handling processor for faster access • Statistics parameters are now directly read from hardware instead of being calculated by software. • Added support for BlueFlame. BlueFlame is now the default working mode for all packets that have a descriptor which fits into a BF register (currently 256 bytes). Use "BlueFlame" registry key to enable/disable this feature. • Added support for RSS functionality on available processors numbers. Used to be restricted to start at the first processor. • Changed RSS registry defaults to give better out of the box performance • Added a performance UI to tune performance under various scenarios • Added a tool to tune performance under various scenarios
	Ethernet	<ul style="list-style-type: none"> • Added support for multiple TX rings • Added an option to verify that the number of multicast groups used is no higher than the firmware limits • Improved performance in virtualization when using VMQ

5 API Change Log History

Table 21 - API Change Log History

Release	Name	Description
4.95.50000	ND extension for Resource pools	Please refer to MLNX- _VPI_WinOF_User_Manual_v4.95
4.80.50000	RDMA_TRANSPORT_RDMAOE_1	It is an alias to: RDMA_TRANSPORT_RDMAOE
	RDMA_TRANSPORT_RDMAOE_1_25	Added enumerated values
	is_roce(), is_mac_based_roce(), is_ip_based_roce(), is_rro- ce_or_ip_based_roce()	Added new functions
	struct ib_wc_t	p_next was replaced with an anonymous union which contains two fields: p_next and qp_context
4.70	ib_get_port_spl_qp()	Added a new function
	ib_get_mad_inner()	Changed API (one more input parameter was added)
	ib_get_mad()	Changed API (one more input parameter was added)
	VERBS_MINOR_VER	Increased its value, 0x000a -> 0x000c
	UNBOUND_PORT_NUM	Added a new macro
4.60	IB_MOD_QP_CHANGE_COUNTER_INDEX	Added a new macro
	struct ib_qp_mod_t	Added the field state.rtr.counter_index
4.40	VERBS_MINOR_VER	Increased its value, 0x0009 -> 0x000a
	enum eth_link_speeds	Added enumerated values
	struct ib_port_attr_t	<ul style="list-style-type: none"> The mtu field was separated into two fields: <ul style="list-style-type: none"> max_mtu (maximum MTU sup- ported by the port) active_mtu (actual MTU which the port is configured with) Added the eth_link_speed field
	WR_SEND_INV	Added enumerated values
	struct ib_send_wr_t	The type of invalidate_rkey was changed from net32_t -> ib_net32_t
	IB_SEND_OPT_SKIP_DOORBELL	Added the send Write flag